Exhibit 3

Harvard Business School

Alain Bonacossa
he/him/his
Research Integrity Officer,
Division of Research and Faculty Development

617.496.6348 (o)
917.292.8375 (m)
abonacossa@hbs.edu

CONFIDENTIAL

Date: June 13, 2023

By Email:
Dr. May R. Berenbaum, maybe@illinois.edu

Dear Dr. Berenbaum,

I am writing to inform you that Harvard Business School (HBS) has reviewed concerns about certain data previously published by Dr. Francesca Gino in the following article:

> Shu, L. L., Mazar, N., Gino, F., Ariely, D., and Bazerman, M. H. (2012). Signing at the beginning makes ethics salient and decreases dishonest self-reports in comparison to signing at the end. *Proceedings of the National Academy of Sciences of the United States of America*, 109, 15197–15200

We have included an appendix to this letter that includes (a) a description of the data anomalies for Study 1 observed in the Open Science Framework (OSF) dataset, (b) a comparison of original data for Study 1 with the data observed in OSF, and (c) an assessment by an independent forensic firm. In summary, a comparison of (a) an original dataset provided by a former research staff member who worked on the study with (b) the dataset posted on OSF revealed a large number of discrepancies. Two types of discrepancies are particularly notable: first, 6 participants' condition assignments differed in the two datasets and, second, 52% of the participants that could be confidently matched had data that were different in the two datasets, with no identified reason for the differences. All but one of these discrepancies favor the hypothesized and reported effects. Moreover, the report pointed to differences in the statistical results for both dependent variables that contradict the published paper.

We thus believe the results reported in Study 1 of the above-referenced paper are invalid due to alteration of the data that affects the significance of the findings. We are informing the article's co-authors, and we recommend amending the article's retraction notice accordingly.

If you wish to discuss this matter further or if you have any questions, please feel free to reach out to me at 617-496-6348 or abonacossa@hbs.edu.

Sincerely,

Alain Bonacossa
Research Integrity Officer

APPENDIX REGARDING
Shu, L. L., Mazar, N., Gino, F., Ariely, D., and Bazerman, M. H. (2012). Signing at the beginning makes ethics salient and decreases dishonest self-reports in comparison to signing at the end. Proceedings of the National Academy of Sciences of the United States of America, 109, 15197–15200.

Data anomalies observed in the OSF dataset (Study 1)

The dataset for Experiment 1 was retrieved from the OSF, where, since 2020, it has been publicly posted (https://osf.io/4b7mu/).

The posted dataset seems to be sorted by two columns, first by a column called "Cond", indicating participants' condition assignment (0 = control; 1 = sign-at-the-top; 2 = sign-at-the-bottom), and then by a column called "P#", indicating a Participant ID number assigned by the experimenter.

For example, this is a screenshot of a few dozen observations from the sign-at-the-top and sign-at-the-bottom condition. One can see that within each condition the data are *almost* perfectly sorted by Participant ID (the first column on the left). However, eight observations are out of order, as highlighted below:

	A	B	C	D	E	F	G	H	I
1	**P#**	**Cond**	**Stude**	**Major**	**CS3**	**Male**	**Age**	**#B**	**$B**
47	35	1	1	Journalism	3	1	19	12	12
48	37	1	1	Economics	4	0	21	9	9
49	40	1	1	Political Science	5	1	29	15	15
50	42	1	1	Political Science	3	0	20	7	7
51	46	1	1	Political Science	4	0	21	12	12
52	49	1	1	English	4	1	21	9	9
53	49	1	1	English	4	1	21	7	7
54	55	1	1	Biology	4	1	21	12	12
55	58	1	1	Environmental Sciences	3	0	20	10	10
56	61	1	1	Nursing	3	0	20	15	15
57	63	1	0	NA		0	22	12	12
58	68	1	1	Business	3	1	20	16	16
59	70	1	1	Chemistry	4	0	21	11	11
60	73	1	1	Chemistry	5	0	20	16	16
61	76	1	1	Chemistry	2	1	19	15	15
62	80	1	1	Nursing	4	0	21	15	15
63	82	1	1	Economics	4	1	21	9	9
64	85	1	1	Psychology	4	0	20	5	5
65	88	1	1	Chemistry	3	0	20	13	13
66	95	1	1	Math Education	3	1	22	13	13
67	51	1	0	NA	0	0	52	4	4
68	12	1	1	Psychology	3	0	20	13	13
69	101	1	0	Business	3	1	20	6	6
70	7	2	0	Political Science	5	1	22	17	17
71	91	2	1	Japanese	2	1	20	17	17
72	52	2	0	NA	5	0	22	8	8
73	5	2	1	Biology/Psychology	2	0	18	16	16
74	8	2	1	Communication Studies	4	0	22	15	15
75	13	2	1	Chemistry	4	0	20	18	18
76	17	2	1	Communications	4	0	21	14	14
77	18	2	1	Communications	4	1	22	13	13
78	22	2	0			0	23	13	13
79	26	2	0			0	47	6	6
80	27	2	1	Mathematics - Sociology	3	1	19	18	18

Participant ID 49 appears twice in the dataset, with identical demographic variables. In addition, Participants 51, 12, 101 are out of order in Condition 1, and Participants 7, 91, and 52 are out of order in Condition 2. This is a red flag because there is no apparent way to sort the data to achieve this ordering. It suggests that observations must have been moved around (or duplicated), manually, perhaps to alter a participant's condition assignment in a way that achieves the desired result.

A deeper dive into the data of these eight participants provides support for this form of data alteration. The figure below shows a "Bee Swarm" plot, which depicts each observation in the dataset, separately for each experimental condition. The plot depicts one of the cheating measures, the amount of money participants claimed in travel expenses. Every "normal", in-sequence observation is represented as a blue

dot, whereas the eight out-of-sequence observations are represented as red X's.

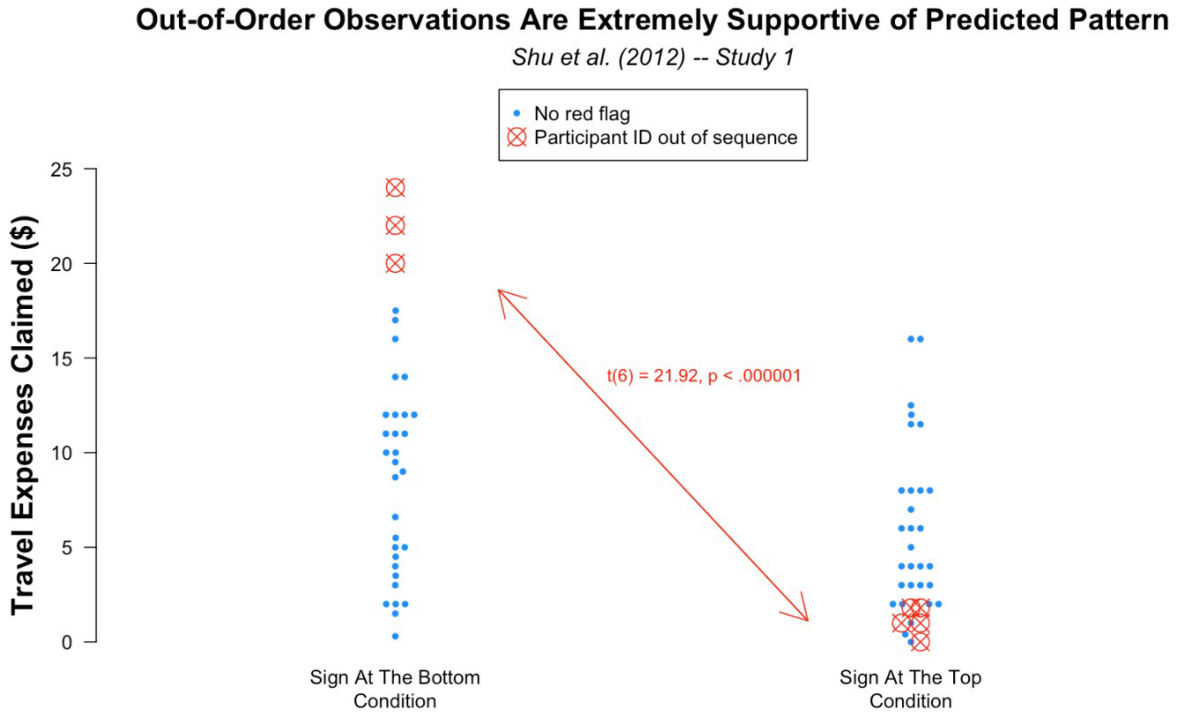

In the sign-at-the-bottom condition, the authors predicted expenses to be high, and the three out-of-sequence observations in this condition are the very highest. In the sign-at-the-top condition, the authors predicted expenses to be low, and the five out-of-sequence observations in this condition were all among the very lowest. As shown in the plot, the condition difference between just these eight observations on this dependent variable is very highly significant; it would occur by chance less than 1 in a million times.

A similar effect emerges when analyzing the other dependent variable, the overreporting of the number of math puzzles solved. The five out-of-sequence observations in the sign-at-the-top condition, predicted to be low, are all equal to zero, the lowest value observed in the dataset. The three out-of-sequence observations in the sign-at-the-bottom condition, predicted to be high, were all greater than zero: 2, 6, and 7. The condition difference between these eight observations on this dependent variable was again highly significant, even with few observations: $t(6) = 4.48$, $p = .004$.

In sum, there are eight observations that are out of order in this dataset, and it appears that no sorting function can account for their placement. This suggests that these eight observations may have been altered to produce the desired effect. Supporting that contention, those eight observations play a sizable role in producing the published effect in Study 1, as all eight observations have values on the dependent variables that are extremely consistent with the authors' hypothesis.

Comparison of observations between the original dataset and OSF dataset (Study 1)

When the anomalous observations were removed from the dataset, the mean score on Travel Expenses of the "Signature at Top" group increased from 5.3 to 6.0, and the mean score of the "Signature at Bottom" group decreased from 9.6 to 8.4. The adjustment reduced the difference between the two groups in a direction opposite to that of the authors' hypothesis.

Assessment by an Independent Forensic Firm

Executive Summary.

Within the data files reviewed there appear to be multiple discrepancies in various areas of the original data source(s) ("***Excel Files***") and public repository data associated with the 2012 PNAS Paper ("***OSF data***"). The discrepancies include alterations of data included and reported across all three "treatment areas": Condition 0 "No signature", Condition 1 "Signature at the top", and Condition 2 "Signature at the bottom". Furthermore, assessment areas of both "math puzzles reported" as well as "claimed deductions" for the three treatment conditions appear to be modified with directionality (e.g., comparative alterations appear to align with described theorized and resultant published behavioral modifications). Statistical analysis of the data uploaded to the OSF site is consistent with the results reported in the 2012 PNAS Paper. Utilizing the same calculations for the Excel Files demonstrates that a) outcomes appear contrary to reported study effects and b) often have lower (or no) statistical significance.

ANALYSIS AND OBSERVATIONS

Data Analysis.

Publicly available Study 1 data (***OSF data*** from the site "Reducing Dishonesty – Replication(s)" https://osf.io/2ehzt/) was compared to reported original data (***Excel files data***). The ***Excel files*** were provided by the research assistant (**"RA"**) aiding with the study.

Observations and Summary.

i. **Comparing the *OSF data* (published) to the *Excel files* (original):**
   - Six participant IDs are present in both the ***OSF data*** and the ***Excel files*** but under different conditions.
   - **61% of all reported response entries** that were successfully matched between the ***Excel files data*** and the ***OSF data,*** **have been modified** in the ***OSF data*** as compared to the original data.
     - **9% of the modified reported responses above** contained miscalculations that were apparently corrected by the authors in the ***OSF data***.
     - **52% of reported responses contained entries that were modified** without apparent cause. Within each participant, the extent of modification is explored below.

ii. **Comparing the *OSF data* (published) to the *Excel files* (original), one by one comparison of results:**

To aid data visualization and summarize the evaluation of the two data sets we provide a heatmap of the difference in scores by subtracting the Excel files score from the OSF score. Here, we are only comparing the values that were matched within specific

conditions and excluded the 6 survey responses that were previously identified as having modified conditions. When inspecting the single survey responses entries, trends of modifications may not be readily discernable. However, when the additional calculations, and intermediate steps before plotting and statistical analysis are computed, a certain trend becomes apparent.

Discussion, Condition 0

When considering Condition 0, the *no signature* condition, 16/31 survey responses (approximately 52%) appear to have been modified by increasing the rate of over-reporting of solved matrixes **1**, and/or of participants who 'cheated' on the matrix task **2**, and/or of deductions claimed **3** (see Condition 0 data table below). In some cases, all three reported outcome data areas were modified per participant.

Condition 0 data table, snapshot

Participant ID	1 delta OverReport	2 delta CheatedOnMatrixTax	3 delta Deductions
2	0	0	2
3	4	1	8
10	0	0	1
13	2	0	4
15	0	0	0
19	0	0	0
21	0	0	0
24	0	0	2
29	0	0	2
32	0	0	6
33	0	0	0
39	3	1	4
41	1	0	0
45	0	0	0
47	2	1	4
50	0	0	0
53	0	0	1
56	4	1	0
59	0	0	0
62	0	0	1
65	3	1	2
67	0	0	0
71	1	0	0
74	0	0	0
79	0	0	0
81	0	0	0
83	6	1	0
87	0	0	0
89	0	0	0
93	0	0	0
96	0	0	0

MCG Discussion, Condition 1.

When considering Condition 1, the *signature at the top* condition, 7/31 survey responses (approximately 26%) appear to have been modified by *decreasing* the rate of over-reporting of deductions claimed 3, and this appeared to be the case for the majority of the outcome data modified. One participant had two of the outcomes modified by decreasing the rate of over-reporting of solved matrixes 1, and/or of participants who 'cheated' on the matrix task 2 (Participant 63). One exception is the trends for Condition 1 data are represented by participant 1, who appears to have under-reported the number of matrixes solved. In this instance, the survey response was adjusted to report the correct number of matrixes that had been claimed (e.g., the participant apparently underreported their performance, and this was modified to align the data as reported performance = actual performance).

Condition 1 data table, snapshot

Participant ID	1 delta OverReport	2 delta CheatedOnMatrix Tax	3 delta Deductions
1	1	0	0
4	0	0	0
6	0	0	0
9	0	0	0
11	0	0	-2
14	0	0	0
16	0	0	0
20	0	0	0
23	0	0	-5
25	0	0	0
28	0	0	-8
31	0	0	0
35	0	0	0
37	0	0	0
40	0	0	-16
42	0	0	0
46	0	0	0
49	0	0	0
55	0	0	0
58	0	0	-4
61	0	0	0
63	-6	-1	0
68	0	0	-2
70	0	0	0
73	0	0	0
76	0	0	0
80	0	0	0
82	0	0	0

85	0	0	0
88	0	0	0
95	0	0	0

MCG Discussion, Condition 2.

When considering Condition 2, the *signature at the bottom* condition, 13/30 survey responses (approximately 43%) appear to have been modified by *increasing* the rate of over-reporting of deductions claimed **3**, and this appeared to be the case for the majority of the outcome data modified. Three participants had all three outcomes (**1, 2** and **3**) modified (Participants 34, 38 and 97) and one participant (8) had both over-reporting of solved matrixes **1** and deductions claimed **3** modified.

Condition 2 data table, snapshot

Participant ID	**1** delta OverReport	**2** delta CheatedOnMatrix Tax	**3** delta Deductions
5	0	0	0
8	1	0	10
13	0	0	0
17	0	0	0
18	0	0	0
22	0	0	0
26	0	0	0
27	0	0	0
30	0	0	0
34	4	1	8
36	0	0	0
38	5	1	7
43	0	0	6
44	0	0	4
48	0	0	0
54	0	0	0
57	0	0	0
60	0	0	1
66	0	0	4
69	0	0	3
72	0	0	0
75	0	0	6
77	0	0	2
78	0	0	0
84	0	0	0
86	0	0	8
90	0	0	0
92	0	0	2
94	0	0	0
97	4	1	8

iii. **Comparing the *OSF data* (published) to the *16-Jul data* (original) quantitative trends and outcomes:**

The replication of the statistical assessment of the data relative to Study 1/Experiment 1 often shows lower (or no) statistical significance between various stated outcome groups when comparing results obtained under the three conditions (for example, see Table 1 and section **d.** below). It also shows different trends of means between the ***OSF data*** and the data from the ***Excel files*** when considering the percentage of cheating participants, the number of math puzzles over reported, the number of deductions claimed and the significance between conditions (for example, see section **d.** below for discussion).

a. Evaluation of participants who cheated on matrix tax, Figure 1. and Table 1.

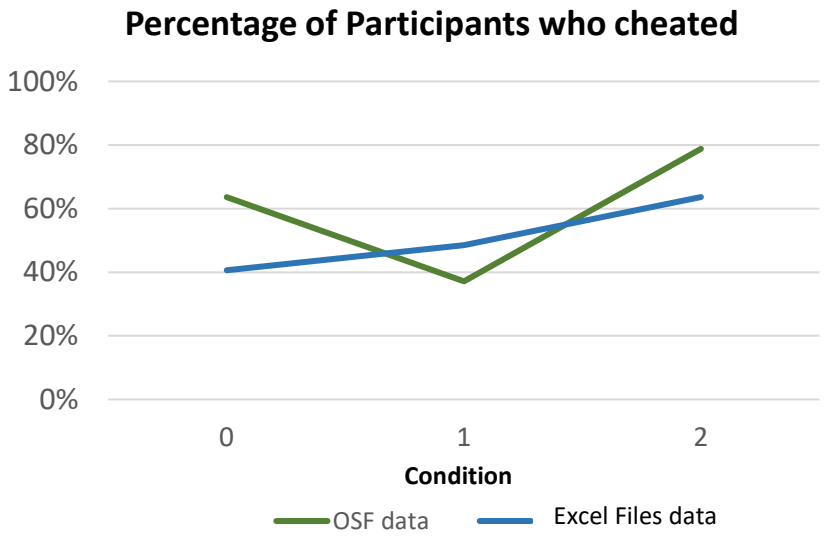

Figure 1. From *Summary Analysis Statistics* sheet. Differences across all conditions in the reported percentage of participants who cheated when comparing the published ***OSF data*** to the original ***Excel files data***. "No Signature" = 0, "Signature at Top" = 1 "Signature at Bottom" = 2.

Table 1. Participants and χ^2 snapshot of "participants who cheated"

		χ2		CheatedOnMatrixTax		
	N of participants	value	p	0	1	2
OSF data	101	12.58	0.002	64%	37%	79%
Excel files data	98	3.57	0.168	41%	48%	64%

While the published ***OSF data*** reported that the percentage of participants who cheated differed significantly across conditions, and significantly less if asked to sign at the top (Figure 1, **GREEN** trend line) the ***Excel files data*** (Figure 1, **BLUE** trend line) do not appear to demonstrate similar trends with respect to the inter-condition

differences. As an example, compare the ***OSF data*** line to the ***Excel files data*** line between Conditions 0 and 1; the directionality is opposite. Additionally, the *overall* inter-condition trends are different (e.g., Condition 1< Condition 0 *and* 2 in ***OSF data*** where Condition 0< Condition 1< Condition 2 in ***Excel files data***). Specifically, the 2012 PNAS paper ***OSF data*** described that fewer participants cheated in the signature-at-the-top condition (Condition 1, 37%) than in the signature-at-the-bottom (Condition 2) and no-signature (Condition 0) conditions (79 and 64%, respectively), where the ***Excel files data*** appear to demonstrate that the no-signature condition (Condition 0) had the fewest participants cheating (41%) than either the signature-at-the-top (Condition 1, 48%) or signature-at-the-bottom (Condition 2, 64%) conditions. Furthermore, the significance between the resultant data per conditions appear reduced or absent in the ***Excel files data*** in comparison to the published ***OSF data.*** This pattern of differences when comparing the ***OSF data*** to the ***Excel files data*** in effect size between conditions and significance in reported results is found in other areas of the datasets as well. For example, in both Math puzzles overreported and Claimed deductions, the effects between conditions and significance of reported results are altered, and in many cases reduced or absent ***Excel files data*** when compared to ***OSF data***.

b. Number of math puzzles over reported, Figure 2 and Table 2.

In the 'OverReporting Math puzzles per Condition' the data trends are altered in that there appears to be more Math puzzles overreported for the "No signature" and the "signature at the bottom" conditions and less puzzles over reported for the "signature at the top" condition in the ***OSF data*** then the original ***Excel files data***. Additionally, the overall values for all conditions are different between the two data sets.

Table 2. Math puzzles overreported snapshot

	M			SD		
Condition	0	1	2	0	1	2
OSF data	2.52	0.77	3.94	3.12	1.44	4.07
Excel Files data	1.41	1.36	3.18	2.99	2.18	4.25

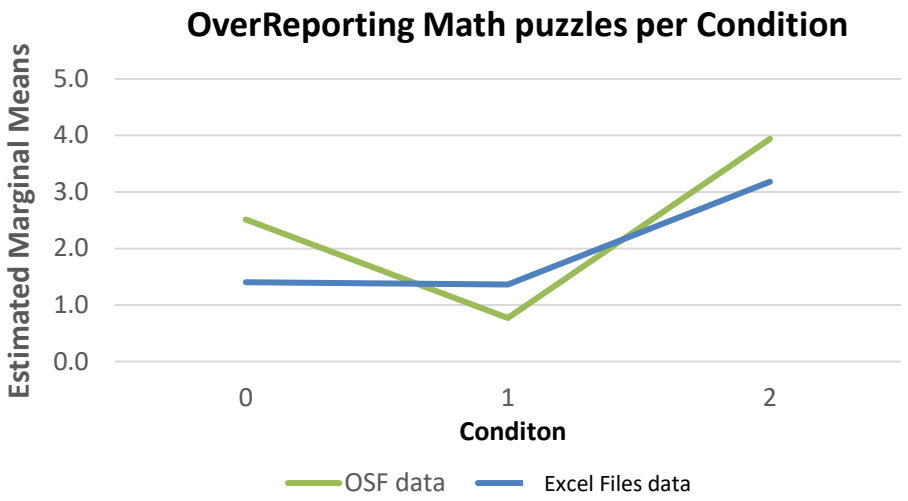

Figure 2. From *Summary Analysis Statistics* sheet. Differences across all conditions in the estimated marginal means of participants overreporting on math puzzles when comparing the published ***OSF data*** to the original ***Excel files data***. "No Signature" = 0, "Signature at Top" = 1 "Signature at Bottom" = 2.

Figure 3. Example of Relative effects on Published Figure 1 Data. When comparing the ***OSF data*** to the ***Excel files data*** in the calculated number of math puzzles solved (e.g., as a relative measure of honesty when comparing described reported data vs actual data), the trends for each Condition appear opposite; both in absolute values (inter-data set differences, for example compare **SOLID GREEN COLUMNS** to **SOLID BLUE COLUMNS**) and magnitude (intra-data set differences, for example compare **GREEN dashed trend line** to **BLUE**).

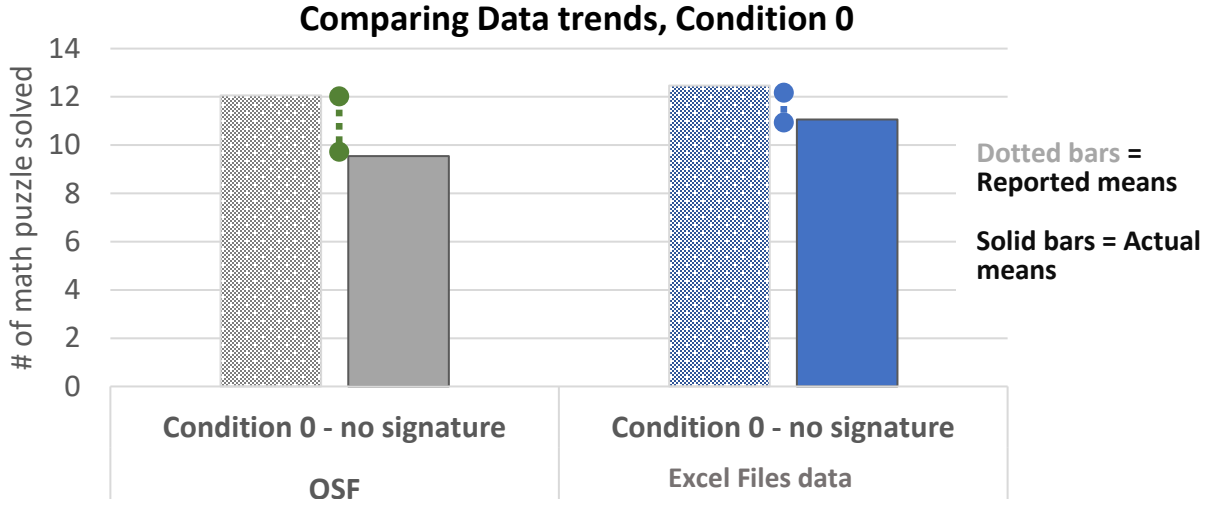

Figure 3. From *Summary Analysis Data* sheet. **Condition 0** "No Signature" data for ***OSF (left graphs)*** vs ***Excel files (right graphs)*** data sets. Inter-data set differences: there are overall *a smaller* number of solved puzzle data included in the **OSF** than in the ***Excel files data*** (# puzzles solved **OSF** < ***Excel files***

data *for both* ***Actual*** *and* ***Reported*** *data*). Intra-data set differences: the difference between the Actual and Reported puzzles solved for the ***OSF data*** is *greater than* the difference between the Actual and Reported puzzles solved for the ***Excel files data*** (△ "Actual vs Reported" ***OSF data*** > △ "Actual vs Reported" ***Excel files data.***)

c. Number of deductions requested, Figure 4 and Table 4.

In the 'Claimed deductions per Condition', in addition to the overall values for all conditions being different (**Table 4.**, Means [M] and Standard deviations [S]), the trend for Conditions 1 ("Signature at the Top") and 2 ("Signature at the bottom") appear to be opposite in the original ***Excel files data*** compared to the ***OSF data*** (**Figure 4.**)

Table 4. Claimed Deductions snapshot

	M			SD		
Condition	0	1	2	0	1	2
OSF data	8.45	5.27	9.62	5.92	4.43	6.20
Excel files data	8.12	7.93	5.90	6.26	6.95	5.12

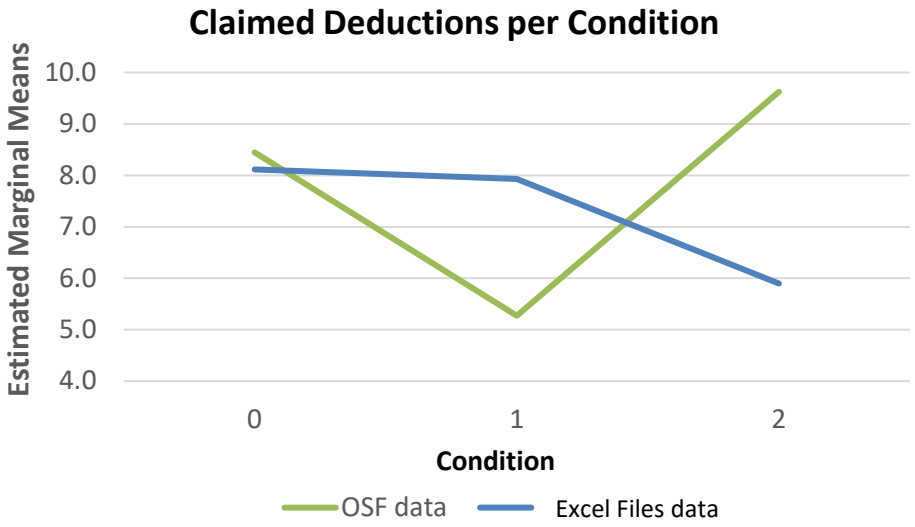

Figure 4. 'Claimed deductions' are higher for Condition 1 in the original ***Excel files data*** vs the ***OSF data*** and lower for Condition 2 in the ***Excel files data*** vs the ***OSF data***.

d. Other statistical inconsistencies:

Other instances of discrepancies between statistical results obtained with the ***OSF data*** and ***Excel data*** can be found in the table below:

Table 4. Additional statistical results

	Math puzzles overreported						Claimed deductions					
	between subjects effects			Between conditions effects			between subjects effects			Between conditions effects		
				p (LSD post hoc)						p (LSD post hoc)		
	F	η^2	p	0-1	1-2	0-2	F	η^2	p	0-1	1-2	0-2
Published data	9.21	0.16	<0.001	<0.05	<0.001	<0.07	5.63	0.10	<0.01	<0.05	<0.01	0.39
OSF data	9.21	0.16	0.0002	0.02	0.00004	0.06	5.63	0.10	0.005	0.02	0.002	0.39
Excel files data	3.33	0.07	0.04	0.96	0.03	0.03	1.31	0.03	0.27	0.90	0.18	0.15

As outlined above, the data as reported in the 2012 PNAS paper appear to align with the ***OSF data*** analyzed using SPSS software. However, when considering statistical significance, the F test analysis' p-value for both 'math puzzles overreported' and 'claimed deductions', appears to be higher in the ***Excel files data*** than in the ***OSF data*** (2012 PNAS), and higher than 0.05 for both (see "F" in Table 4, p column for "math puzzles over reported", 0.04 (***Excel files***) compared to 0.0002 (**OSF**) and, p column for "claimed deductions", 0.27 (***Excel files***) compared to 0.005 (**OSF**)).

Comparing calculated η^2, for both data sets, the OSF η^2 values align with the 2012 PNAS published data for both 'math puzzles over reported' and 'claimed deductions (see "F" in Table 4. η^2 columns for both 2012 PNAS and OSF; aligning values are 0.16 and 0.10, respectively). When the same η^2 calculations are completed for the ***Excel files data***, the resultant values appear decreased compared to their 2012 PNAS and OSF counterparts (see "F" in Table 4. η^2 column for "math puzzles over reported", 0.07 (***Excel files***) compared to 0.16 (**OSF**) and, η^2 column for "claimed deductions", 0.03 (***Excel files***) compared to 0.10 (**OSF**)).

To compute the significance between conditions, all post-hoc Anova significance algorithms were tested. A match was found with the published results when using Fisher's Least Significant Difference (LSD) Test post-hoc on ***OSF data***. The same algorithm on the Excel files data shows lower significance. As evident from Table 4. the ***Excel files data*** show lower significance of the difference of effect between conditions (compare LSD post hoc analysis of ***Excel files*** dataset to ***OSF data*** sets, respectively). Most importantly, the p values across conditions appear to be consistently above 0.05 for "Claimed deductions" as well as for the difference between conditions 0 and 1, "no signature" and "signature at the top condition" for "math puzzles over reported".

Summary of the Study 1/Experiment 1 *OSF data (published)* to the *Excel files data (original)* data set comparisons:

The trend of the data alterations throughout the respective areas of the ***OSF data*** appears to align with the authors theorized projections published for Condition 1; that signing before self-reporting data appeared to be a more effective reinforcement of "honest behavior" than signing afterward (Condition 2) or not signing at all (Condition 0).

- "math puzzles reported" as well as "claimed deductions" for the three treatment conditions appear to be modified with directionality (e.g., comparative alterations appear to align with described theorized and resultant published behavioral modifications).
- The Excel Files results show that a) outcomes appear contrary to published study effects and b) often have lower (or no) statistical significance.